arXiv:1703.05863v2 [cs.CG] 18 Apr 2017

Packing Short Plane Spanning Graphs in Complete Geometric Graphs*

Oswin Aichholzer† Thomas Hackl† Matias Korman‡
Alexander Pilz§ Günter Rote¶ André van Renssen‖, **
Marcel Roeloffzen‖,** Birgit Vogtenhuber†

April 19, 2017

Abstract

Given a set of points in the plane, we want to establish a connection network between these points that consists of several disjoint layers. Motivated by sensor networks, we want that each layer is spanning and plane, and that no edge is very long (when compared to the minimum length needed to obtain a spanning graph).

We consider two different approaches: first we show an almost optimal centralized approach to extract two graphs. Then we show a constant factor approximation for a distributed model in which each point can compute its adjacencies using only local information. In both cases the obtained layers are plane.

1 Introduction

Given a set S of n points in the plane and an integer k, we are interested in finding k edge-disjoint non-crossing spanning graphs $H_1, H_2, \ldots, H_k$ on S such that the length $\mathsf{BE}(H_1 \cup H_2 \cup \cdots \cup H_k)$ of the *bottleneck edge* (the longest edge which is used) is as short as possible. Each subgraph H_i is referred to as a *layer* of G. We require each layer to be non-crossing, but edges from different layers are allowed to cross each other. For $k = 1$, the minimum spanning tree $\mathsf{MST}(S)$ solves the problem: its longest edge $\mathsf{BE}(\mathsf{MST}(S))$ is a lower bound

*This research was initiated during the 10th European Research Week on Geometric Graphs (GGWeek 2013), Illgau, Switzerland.

†Institute of Software Technology, Graz University of Technology, Austria. {oaich|thackl|bvogt}@ist.tugraz.at

‡Tohoku University, Sendai, Japan. mati@dais.is.tohoku.ac.jp

§Department of Computer Science, ETH Zürich, Switzerland. alexander.pilz@inf.ethz.ch

¶Institut für Informatik, Freie Universität Berlin, Germany. rote@inf.fu-berlin.de

‖National Institute of Informatics, Tokyo, Japan. {andre|marcel}@nii.ac.jp

**JST, ERATO, Kawarabayashi Large Graph Project.

on the bottleneck edge of any spanning subgraph, and it is non-crossing. For larger k, we take $\mathsf{BE}(\mathsf{MST}(S))$ as the yardstick and measure the solution quality in terms of $\mathsf{BE}(\mathsf{MST}(S))$ and k.

The particular variation that we consider comes motivated from the field of sensor networks. Imagine one wants to construct a network so that afterwards communication between sensors is possible. One of the most important requirements for such a network is that we can send messages through it easily. Ideally, we want a method that – given the source, destination, information on the current position (and possibly $O(1)$ additional information) – computes the next node to visit in order to reach our destination.

One of the most famous such methods is called *face routing* [7], which guarantees the delivery under the above constraints provided that the underlying graph is plane. Indeed, when considering local routing algorithms in the literature that are guaranteed to succeed, most route deterministically on a plane spanning subgraph of the underlying graph where the plane subgraph can be computed locally. Even though there exist routing strategies for non-plane graphs, in most cases they route through a plane subgraph. (For example, Bose *et al.* [2] showed how to locally identify the edges of the Gabriel graph from the unit disk graph). Extending these algorithms for non-plane graphs is a long-standing open problem in the field.

It seems counterintuitive that having additional edges cannot help in the delivery of messages. In this paper, we provide a different way to avoid this obstacle. Rather than limiting considerations to one plane graph, we aim to construct several disjoint plane spanning graphs. If we split all the messages among the different layers (and route through each layer with routing strategies that work on plane graphs) we can potentially spread the load among a larger number of edges. Another important feature to consider when creating networks is energy consumption. The required energy for sending a message from one point to the next increases with the distance between the two points (usually with the third or fourth power) [4]. Since we want to avoid high energy consumption at one particular node, it is desirable to apply the bottleneck criterion and to minimize the longest edge [6].

Previous Work. This problem falls into the family of *graph packing* problems, where we are given a graph $G = (V, E)$ and a family $\mathcal{F}$ of subgraphs of G. The aim is to pack as many pairwise disjoint subgraphs $H_1 = (V, E_1), H_2 = (V, E_2), \ldots$ as possible into G.

A related problem is the *decomposition* of G. In this case, we also look for disjoint subgraphs but require that $\bigcup_i E_i = E$. For example, there are known characterizations of when we can decompose the complete graph of n points into paths [9] and stars [8]. Dor and Tarsi [3] showed that to determine whether we can decompose a graph G into subgraphs isomorphic to a given graph H is NP-complete. Concerning graph packing, Aichholzer *et al.* [1] showed that the complete graph on any set S of n points contains $\Omega(\sqrt{n})$ disjoint plane spanning trees. This bound has been improved to $\lfloor n/3 \rfloor$ by Garcia [5].

In our case, the graph G consists of the complete graph on a given point set S, and $\mathcal{F}$ consists of all plane spanning graphs of G. In addition to maximizing the number of layers, we are interested in minimizing a geometric constraint (Euclidean length of the longest edge among the selected graphs of $\mathcal{F}$). To the best of our knowledge, this is the first packing problem of such type.

Results. We give two different approaches to solve the problem. In Section 2 we give a construction for two spanning trees, i.e., $k = 2$. This construction is centralized in a classic model that assumes that the positions of all points are known and computed in a single place. Our construction creates two trees and guarantees that all edges (except possibly one) have length at most $2\mathsf{BE}(\mathsf{MST}(S))$. The remaining edge has length at most $3\mathsf{BE}(\mathsf{MST}(S))$. We complement this construction with a matching worst-case lower bound.

Following the spirit of sensor networks, in Section 3 we use a different approach to construct k disjoint plane graphs (not necessarily trees). The construction works for any $k \leq n/12$ in an almost local fashion. The only global information that is needed is β: $\mathsf{BE}(\mathsf{MST}(S))$ or some upper bound on it. Each point of S can compute its adjacencies by only looking at nearby points, namely, those at distance $O(k\beta)$.

A simple adversary argument shows that it is impossible to construct spanning networks locally without knowing $\mathsf{BE}(\mathsf{MST}(S))$ (or an upper bound). The lower bound of Section 2 shows that a neighborhood of radius $\Omega(k\mathsf{BE}(\mathsf{MST}(S)))$ may be needed for the network, so we conclude that our construction is asymptotically optimal in terms of the neighborhood.

For simplicity, throughout the paper we make the usual general position assumption that no three points are collinear. Without this assumption, it might be impossible to obtain more than a single plane layer (for example, when all points lie on a line). Note however, that if collinear and partially overlapping edges are considered as non-crossing, our algorithms do not require the point set to be in general position.

2 Centralized Construction

In this section we look for a centralized algorithm to construct two layers. We start with some properties on the minimum spanning tree of a set of points.

Lemma 1. *Let S be a set of points in the plane and let uv and vw be two edges of $\mathsf{MST}(S)$. Then the triangle uvw does not contain any other point of S.*

Proof. Observe that, as v is adjacent to both u and w in $\mathsf{MST}(S)$, uw is the longest edge of the triangle uvw (otherwise one could locally shorten $\mathsf{MST}(S)$).

Suppose for the sake of contradiction that there is a point $p \in S$ in the interior of uvw. We split uvw into two sub-triangles by the line ℓ through v perpendicular to the supporting line of u and w. Let Δ_u be the sub-triangle that has u as a vertex, and assume w.l.o.g. that p lies in Δ_u. Note that the edge uv is the hypotenuse of the right-angled triangle Δ_u and hence $\max\{|pu|, |pv|\} < |uv|$.

Consider the paths in $\mathsf{MST}(S)$ from p to u and v, respectively. Since $\mathsf{MST}(S)$ is a tree, one of the two paths must use the edge uv (as otherwise there would be a cycle). Suppose first that this edge is used in the path to u. By removing the edge uv and adding the edge pu to $\mathsf{MST}(S)$ we would obtain a connected (not necessarily plane) tree whose overall weight is smaller, a contradiction. If the edge used is in the path to v, the addition of edge pv yields a similar contradiction. □

Lemma 2. *Let S be a set of points in the plane. Let $v \in S$ be a point of degree $k \geq 3$ in $\mathsf{MST}(S)$, with $\{v_0, \ldots, v_{k-1}\}$ being the neighbors of v in $\mathsf{MST}(S)$ in counterclockwise order around v. Then for every triple (v_{i-1}, v_i, v_{i+1}) (indices modulo k), the neighbors of v_i in $\mathsf{MST}(S)$ are inside the wedge W_i that is bounded by the rays vv_{i-1} and vv_{i+1} and contains the edge vv_i.*

Proof. Let $u \in S \setminus \{v\}$ be a neighbor of v_i in $\mathsf{MST}(S)$, and assume for the sake of contradiction that u is not in W_i. Then the edge v_iu intersects the boundary of W_i and, hence, one of the rays starting at v and going through v_{i-1} and v_{i+1}, respectively. Assume without loss of generality that v_iu intersects the ray from v through v_{i+1}. As $\mathsf{MST}(S)$ is plane, the edge v_iu does not intersect the edge vv_{i+1}. Hence, the triangle (v, v_i, u) contains the point v_{i+1} in its interior. As the path vv_iu is a subgraph of $\mathsf{MST}(S)$, this is a contradiction to Lemma 1. □

We denote by $\mathsf{MST}^2(S)$ the square of $\mathsf{MST}(S)$, the graph connecting all pairs of points of S that are at distance at most 2 in $\mathsf{MST}(S)$. We call the edges of $\mathsf{MST}(S)$ *short* edges and all remaining edges of $\mathsf{MST}^2(S)$ *long* edges. For every long edge uw, the points u and w have a unique common neighbor v in $\mathsf{MST}(S)$, which we call the *witness* of uw. We define the *wedge* of uw to be the area that is bounded by the rays vu and vw and contains the segment uw. Next we state a simple fact on crossings of the edges in $\mathsf{MST}^2(S)$.

Lemma 3. *Let S be a set of points in the plane. Two edges e and f of $\mathsf{MST}^2(S)$ cross if and only if one of the following two conditions is fulfilled:*

1. *At least one of $\{e, f\}$ is a long edge with witness v and wedge W, and the other edge has v as an endpoint and lies inside W.*
2. *Both of $\{e, f\}$ are long edges with the same witness v, and their wedges are intersecting but none is contained in the other.*

Proof. Clearly, if both e and f are short edges, i.e., edges of $\mathsf{MST}(S)$, then they do not cross. Without loss of generality assume that $f = uw$ is a long edge with witness v and wedge W. Every edge $e = vz$ of $\mathsf{MST}^2(S)$, $z \in S \setminus \{u, v, w\}$ that lies inside W either crosses f or has z inside the triangle $\Delta = (u, v, w)$. The latter is a contradiction to Lemma 1. Obviously, f is neither crossed by any edge incident to u or w, nor crossed by any edge incident to v but not lying inside W.

It remains to prove that every long edge $e = xz$ of $\mathsf{MST}^2(S)$, $x, z \in S \setminus \{u, v, w\}$ that crosses f fulfills Condition 2. Note that for e to cross f, either e has an endpoint inside Δ or e is also crossing one edge out of $\{uv, vw\} \in \mathsf{MST}(S)$.

The former is a contradiction to Lemma 1. If e is a short edge (i.e., an edge of $\mathsf{MST}(S)$), then the latter is a contradiction to the planarity of $\mathsf{MST}(S)$. Hence, e is a long edge (with wedge W') and is also crossing one edge g out of $\{uv, vw\} \in \mathsf{MST}(S)$. This also implies that the wedges W and W' intersect in their interiors but none of W, W' is contained in the other. Finally, if e has witness $y \neq v$, then either g has an end point in the triangle xyz or g crosses one edge out of $\{xy, yz\} \in \mathsf{MST}(S)$. Again, the former is a contradiction to Lemma 1 and the latter is a contradiction to the planarity of $\mathsf{MST}(S)$. Hence the witness of e must be v, which completes the proof. □

2.1 Constructing two almost disjoint layers

With the above observations we can proceed to show a construction that almost works for two layers. To this end we consider the minimum spanning tree $\mathsf{MST}(S)$ to be rooted at a leaf r. For any $v \in S$, we define its *level* $\ell(v)$ as its distance to r in $\mathsf{MST}(S)$. That is, $\ell(v) = 0$ if and only if $v = r$. Likewise, $\ell(v) = 1$ if and only if v is adjacent to r etc.

For any $v \in S \setminus \{r\}$, we define its *parent* $p(v)$ as the first vertex traversed in the unique shortest path from v to r in $\mathsf{MST}(S)$. Similarly, we define its *grandparent* $g(v)$ as $g(v) = p(p(v))$ if $\ell(v) \geq 2$ and as $g(v) = r$ otherwise (i.e., $g(v) = p(v) = r$ if $\ell(v) = 1$). Each vertex q for which $v = p(q)$ is called a child of v.

Construction 1. *Let S be a set of points in the plane and let $\mathsf{MST}(S)$ be rooted at an arbitrary leaf $r \in S$. We construct two graphs $R = G(S, E_R)$ and $B = G(S, E_B)$ as follows: For any vertex $v_o \in S$ whose level is odd, we add the edge $v_o p(v_o)$ to E_R and the edge $v_o g(v_o)$ to E_B. For any vertex $v_e \in S \setminus \{r\}$ whose level is even, we add the edge $v_e g(v_e)$ to E_R and the edge $v_e p(v_e)$ to E_B.*

For simplicity we say that the edges of $R = G(S, E_R)$ are colored red and the edges of $B = G(S, E_B)$ are colored blue. An edge in both graphs is called red-blue.

Theorem 1. *Let MST(S) be rooted at r. The two graphs $R = G(S, E_R)$ and $B = G(S, E_B)$ from Construction 1 fulfill the following properties:*

1. *Both R and B are plane spanning trees.*
2. $\max\{\mathsf{BE}(R), \mathsf{BE}(B)\} \leq 2\mathsf{BE}(\mathsf{MST}(S))$.
3. *$E_R \cap E_B = \{rs\}$, with $r = p(s)$, that is, $|E_R \cap E_B| = 1$.*

Proof. Recall from Construction 1 that r is a leaf of $\mathsf{MST}(S)$. Hence r has a unique neighbor s in $\mathsf{MST}(S)$ and we have $r = p(s) = g(s)$ and $\ell(s) = 1$. Let $S_o \subset S \setminus \{s\}$ be all $v_o \in S$ whose level $\ell(v_o)$ is odd. Likewise, let $S_e \subset S \setminus \{r\}$ be all $v_e \in S$ whose level $\ell(v_e)$ is even. By construction, E_R contains all the edges from odd-leveled nodes to their parents, those from even-leveled nodes to their

grandparents and rs. More formally,

$$E_R = \bigcup_{v_o \in S_o} \{v_o p(v_o)\} \cup \bigcup_{v_e \in S_e} \{v_e g(v_e)\} \cup \{rs\}.$$

Similarly, E_B contains edges from odd-leveled nodes to their grandparents, those from even-leveled nodes to their parents and rs, that is

$$E_B = \bigcup_{v_o \in S_o} \{v_o g(v_o)\} \cup \bigcup_{v_e \in S_e} \{v_e p(v_e)\} \cup \{rs\}.$$

Thus, the edge rs is the only shared edge between the sets E_R and E_B, as stated in Property 3.

As E_R and E_B are subsets of the edge set of $\mathsf{MST}^2(S)$, the vertices of every edge in E_R and E_B have link distance at most 2 in $\mathsf{MST}(S)$, and the bound on $\max\{\mathsf{BE}(R), \mathsf{BE}(B)\}$ stated in Property 2 follows.

Further, both R and B are spanning trees, that is, connected and cycle-free graphs, as each vertex except r is connected either to its parent or grandparent in $\mathsf{MST}(S)$. To prove Property 1, it remains to show that both trees are plane.

Assume for the sake of contradiction that an edge f is crossed by an edge e of the same color. Recall that all edges of E_R and E_B are edges of $\mathsf{MST}^2(S)$ whose endpoints have different levels. By Lemma 3, at least one of $\{e, f\}$ has to be a long edge. Without loss of generality let $f = uw$ be a long edge and let v be the witness of f with $\ell(u) = \ell(v) - 1 = \ell(w) - 2$. First note that v cannot be an endpoint of e due to its level. That is, uv is not crossing f (due to the common endpoint) and all other edges incident to v in E_R and E_B are either blue if f is red, or red if f is blue. Further, v cannot be the witness of e due to its level. All edges E_R and E_B with witness v have u as one of its endpoints (as for all other edges with witness v in $\mathsf{MST}^2(S)$, both endpoints have the same level). With u as a shared vertex, the edges e and f cannot cross. As e is neither incident to v nor has v as a witness, e crossing f is a contradiction to Lemma 3. This proves Property 1 and concludes the proof. □

The properties of our construction imply a first result stated in the following corollary.

Corollary 2. *For any set S of n points in the plane, there exist two plane spanning trees $R = G(S, E_R)$ and $B = G(S, E_B)$ such that $|E_R \cap E_B| = 1$ and* $\max\{\mathsf{BE}(R), \mathsf{BE}(B)\} \leq 2\mathsf{BE}(\mathsf{MST}(S))$.

Although the construction might seem to generalize to more layers by using edges of $\mathsf{MST}^k(S)$, this is not the case. Already for $k = 4$, there are examples where the corresponding trees are not plane.

2.2 Avoiding the double edge

Construction 1 is almost valid in the sense that only one edge was shared between both trees. In the following we enhance this construction to avoid the shared edge.

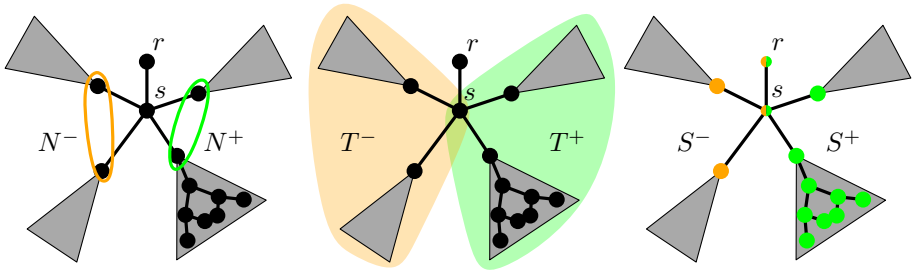

Figure 1: Illustration of the various definitions used in Section 2.2. Grey triangles illustrate further subtrees, one shows interior vertices purely for illustrative purposes.

Let $N^- \subset (S \setminus \{r\})$ be the set of neighbors $v^- \in N^-$ of s in $\mathsf{MST}(S)$ such that the ordered triangle rsv^- is oriented clockwise. Let $N^+ \subset (S \setminus \{r\})$ be the set of neighbors $v^+ \in N^+$ of s in $\mathsf{MST}(S)$ such that the ordered triangle rsv^+ is oriented counterclockwise. Let T^- be the subtree of $\mathsf{MST}(S)$ that is connected to s via the vertices in N^- and let T^+ be the subtree of $\mathsf{MST}(S)$ that is connected to s via the vertices in N^+. Let $S^- \subset S$ consist of r and the set of vertices from T^- and let $S^+ \subset S$ consist of r and the set of vertices from T^+. Observe that $S^- \cap S^+ = \{r, s\}$ (see Figure 1). Let E_R and E_B be sets of red and blue edges as defined in the Construction 1. Then let $E_R^- \subset E_R$ ($E_B^- \subset E_B$) be the subset of edges that have at least one endpoint in $S^- \setminus \{r, s\}$ and let $E_R^+ \subset E_R$ ($E_B^+ \subset E_B$) be the subset of edges that have at least one endpoint in $S^+ \setminus \{r, s\}$. Note that by this definition $E_R = E_R^- \cup E_R^+ \cup \{rs\}$ and $E_B = E_B^- \cup E_B^+ \cup \{rs\}$. With this we define the subgraphs $R^- = G(S^-, E_R^-)$, $R^+ = G(S^+, E_R^+)$, $B^- = G(S^-, E_B^-)$, and $B^+ = G(S^+, E_B^+)$ and prove a useful non-crossing property between these graphs.

Lemma 4. *For any set S of n points in the plane, let $R = G(S, E_R)$ and $B = G(S, E_B)$ be the graphs from Construction 1. Then no edge in E_R^- crosses an edge in E_B^+ and no edge in E_R^+ crosses an edge in E_B^-.*

Proof. Consider any edge $e \in E_R^-$ that is not incident to r. By Lemma 3, such an edge e can be crossed only by an edge incident to at least one vertex of $S^- \setminus \{r, s\}$. Hence, e does not cross any edge of E_B^+.

Assume for the sake of contradiction that there exists an edge $f \in E_B^+$ that crosses an edge $e \in E_R^-$ that is incident to r. By construction, $e = rz$ is a long edge of $\mathsf{MST}^2(S)$ with witness s and wedge W. By Lemma 3, f has to be incident to s, since s cannot be the witness of any blue edges by construction. If f is a short edge, then f is not in W by our definition of S^- and S^+, which is a contradiction to Lemma 3. Hence, let $f = sc$ be a long edge of $\mathsf{MST}^2(S)$ with witness b. Following Lemma 3, the witness b must be s, which is in contradiction to the fact that s cannot be a witness of any blue edge. This concludes the proof that no edge in E_R^- is crossed by an edge in E_B^+. Symmetric arguments prove that no edge in E_R^+ is crossed by an edge in E_B^-. □

With this observation we can now prove that the two spanning trees from Construction 1 actually exist in 4 different color combination variants.

Lemma 5. *Let S be a set of n points in the plane. Let $R = G(S, E_R)$ and $B = G(S, E_B)$ be the graphs from Construction 1 and let $R^- = G(S^-, E_R^-)$, $R^+ = G(S^+, E_R^+)$, $B^- = G(S^-, E_B^-)$, and $B^+ = G(S^+, E_B^+)$ be subgraphs as defined above. Then R and B can be recolored to any of the four versions below, where each version fulfills the properties of Theorem 1.*

(1) $R = G(S, E_R)$ and $B = G(S, E_B)$ (the "original coloring")

(2) $R = G(S, E_B)$ and $B = G(S, E_R)$ (the "inverted coloring")

(3) $R = G(S, E_B^- \cup E_R^+ \cup \{rs\})$ and $B = G(S, E_R^- \cup E_B^+ \cup \{rs\})$ (the "− side inverted coloring")

(4) $R = G(S, E_R^- \cup E_B^+ \cup \{rs\})$ and $B = G(S, E_B^- \cup E_R^+ \cup \{rs\})$ (the "+ side inverted coloring")

Proof. The statement is trivially true for recolorings (1) and (2). It is easy to observe that this corresponds to a simple recoloring. Hence, Properties 2 and 3 of Theorem 1 are also obviously true. By Lemma 4, both R and B are plane for the recolorings (3) and (4) and thus fulfill Property 1 of Theorem 1 as well. □

With these tools we can now show how to construct two disjoint spanning trees. For technical reasons we use two different constructions based on the existence of a vertex v in the minimum spanning tree where no two consecutive adjacent edges span an angle larger than π.

Theorem 3. *Consider a set S of n points in the plane for which the minimum spanning tree $\mathsf{MST}(S)$ has a vertex v where between any two consecutive adjacent edges the angle is smaller than π. Then there exist two plane spanning trees $R = G(S, E_R)$ and $B = G(S, E_B)$ such that $E_R \cap E_B = \emptyset$ and $\max\{\mathsf{BE}(R), \mathsf{BE}(B)\} \leq 2\mathsf{BE}(\mathsf{MST}(S))$.*

Proof. We build the two spanning trees by using the vertex v to decompose the minimum spanning tree into trees where v is a leaf. For each of these subtrees we apply Construction 1 and possibly recolor them in one of the variants from Lemma 5.

Let $S_v = \{v_1, \ldots, v_k\}$ be the set of vertices incident to v in $\mathsf{MST}(S)$, labeled in counterclockwise order as they appear around v. Observe that $k \geq 3$ as otherwise the angle condition from the theorem could not be fulfilled. By Lemma 1, the convex hull of S_v contains no points of S except v. We start by constructing two plane spanning trees of $S_v \cup \{v\}$. The red spanning tree $R_v = G(S_v \cup \{v\}, E_{vR})$ contains all edges incident to v except vv_1, plus the edge v_1v_2 that lies on the convex hull boundary of S_v. The blue spanning tree $B_v = G(S_v \cup \{v\}, E_{vB})$ contains all edges on the convex hull of S_v except v_1v_2, plus the edge vv_1. Observe that R_v and B_v are plane spanning trees, $E_{vR} \cap E_{vB} = \emptyset$, and $\max\{\mathsf{BE}(R_v), \mathsf{BE}(B_v)\} \leq 2\mathsf{BE}(\mathsf{MST}(S_v \cup \{v\}))$.

Next consider a vertex v_i of S_v and let M_i be the maximal subtree of $\mathsf{MST}(S)$ that is connected to v by v_i. Let $S_i \subset S$ be the vertex set of M_i. Note that $M_i = \mathsf{MST}(S_i)$ and that v is a leaf in M_i. Thus, we can use Construction 1 to get spanning trees $R_i = G(S_i, E_{iR})$ and $B_i = G(S_i, E_{iB})$, all rooted at v. The graphs R_i and B_i fulfill the three properties of Theorem 1 and the only edge shared between R_i and B_i is vv_i.

Considering Lemma 3 and the fact that for $i \neq j$ the edges of $E_{iR} \cup E_{iB}$ have no point, except for the root v, in common with $E_{jR} \cup E_{jB}$, it is easy to see that no edge of $E_{iR} \cup E_{iB}$ crosses any edge of $E_{jR} \cup E_{jB}$. In order to join the graphs to two plane spanning trees on S, we adapt them slightly, while keeping the properties of Theorem 1. We first state how we combine the edge sets of the different plane spanning trees to get $R = G(S, E_R)$ and $B = G(S, E_B)$ and then reason why the claim in the theorem is true for this construction.

$$E_R = E_{vR} \cup (E_{3R} \backslash \{vv_3\}) \cup \ldots \cup (E_{kR} \backslash \{vv_k\}) \cup (E^-_{1R} \cup E^+_{1B}) \cup (E^-_{2B} \cup E^+_{2R})$$

$$E_B = E_{vB} \cup (E_{3B} \backslash \{vv_3\}) \cup \ldots \cup (E_{kB} \backslash \{vv_k\}) \cup (E^-_{1B} \cup E^+_{1R}) \cup (E^-_{2R} \cup E^+_{2B})$$

First we add the construction for $S_v \cup \{v\}$ to both edge sets. This is the base to which all other trees will be attached. Then the graphs from the subtrees M_i for $1 \leq i \leq k$ are added to this base. The edges vv_i are already used in R_v or B_v, so we do not add the edges vv_i from both colors E_{iR} and E_{iB}. As both v and v_i are connected to both colors (both spanning trees), the construction stays connected. As we did not add any additional edges the construction obviously stays cycle-free and the edge length bound is maintained.

It remains to argue the planarity of the resulting graphs. By Lemma 3, edges of E_{iR} or E_{iB} that cross any edge of E_{vR} and E_{vB} have to be incident to v. By Lemma 2, only the edges $e_i^- = v_{i-1}v_i$ and $e_i^+ = v_iv_{i+1}$ (indices modulo k) are crossed by edges of $E_{iR} \backslash \{vv_i\}$ and $E_{iB} \backslash \{vv_i\}$.

Using the "original coloring" (see Lemma 5) for R_i and B_i only red edges (edges of $E_{iR} \backslash \{vv_i\}$) cross e_i^- and e_i^+. For any $3 \leq i \leq k$, e_i^- and e_i^+ are blue, i.e., $e_i^-, e_i^+ \in E_{vB}$.

For $i = 1$, the edge e_i^+ is red. In this case, we use the "+ side inverted coloring" (see Lemma 5) for R_i and B_i (and exclude the edge vv_i): $E_{iR} = E^-_{iR} \cup E^+_{iB}$ and $E_{iB} = E^-_{iB} \cup E^+_{iR}$. Using this coloring, all shown properties remain valid (see Lemma 5), all edges from R_i and B_i that cross the blue edge e_i^- remain red, and all edges from R_i and B_i that cross the red edge e_i^+ are now blue.

In a similar manner we fix the case of $i = 2$, where the edge e_i^- is red. We use the "− side inverted coloring" (see Lemma 5) for R_i and B_i (and exclude the edge vv_i): $E_{iR} = E^-_{iB} \cup E^+_{iR}$ and $E_{iB} = E^-_{iR} \cup E^+_{iB}$. Again, all shown properties remain valid (see Lemma 5). All edges from R_i and B_i that cross the red edge e_i^- are now blue, and all edges from R_i and B_i that cross the blue edge e_i^+ remain red.

Hence, with this slightly adapted construction (and coloring), $R = G(S, E_R)$ and $B = G(S, E_B)$ are plane spanning trees that solely use edges of $\mathsf{MST}^2(S)$ and have no edge in common. □

The remaining case considers that for every vertex in $\mathsf{MST}(S)$ there exist two consecutive adjacent edges that span an angle larger than π. In such an $\mathsf{MST}(S)$, every vertex has degree at most three, since the angle between adjacent edges is at least $\pi/3$.

Theorem 4. *Consider a set S of $n \geq 4$ points in the plane for which every vertex in the minimum spanning tree $\mathsf{MST}(S)$ has two consecutive adjacent edges spanning an angle larger than π. Then there exist two plane spanning trees $R = G(S, E_R)$ and $B = G(S, E_B)$ such that $E_R \cap E_B = \emptyset$ and $\max\{\mathsf{BE}(R), \mathsf{BE}(B)\} \leq 3\mathsf{BE}(\mathsf{MST}(S))$ (where at most one edge of $E_R \cup E_B$ is larger than $2\mathsf{BE}(\mathsf{MST}(S))$).*

Proof. As before in Theorem 3 we will use our construction scheme for trees rooted at a leaf for the majority of the points and use a small local construction that avoids double edges. In this case, instead of removing a single vertex v to decompose the tree we use a set of four vertices as follows. We first start at a leaf of $\mathsf{MST}(S)$ to generate a connected graph P with four vertices that is a subgraph of $\mathsf{MST}(S)$. Then we show how to construct R and B for S for the different cases of P in combination with the remainder of $\mathsf{MST}(S)$.

The construction of $P = \{v_3, v_2, v_1, v_0\}$: Let v_3 be a leaf of $\mathsf{MST}(S)$. For the construction of P, we root $\mathsf{MST}(S)$ at v_3. With the *weight* of a vertex we denote the number of vertices in the (sub)tree of which that vertex is a root of (including itself). Hence, the weight of v_3 is n. Further we denote with *big angle* the angle between two successive incident edges at a vertex of $\mathsf{MST}(S)$ that is larger than π.

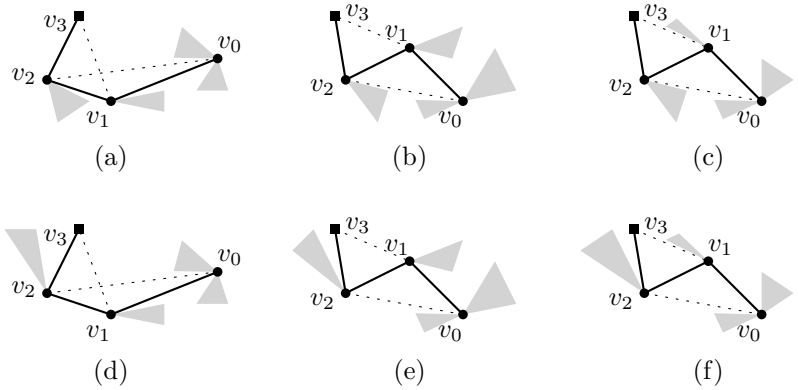

Figure 2: Case 1 for P and the connections to the rest of $\mathsf{MST}(S)$. The gray triangles indicate possible subtrees of $\mathsf{MST}(S)$ and where and how they might be connected. Dotted edges are from $\mathsf{MST}^2(P)$. Note that the subtree with root v_0 can be on either side of the supporting line of v_1v_0 and even on both sides as indicated in the figure.

Let v_2 be the unique child (i.e., $v_3 = p(v_2)$) of v_3 in $\mathsf{MST}(S)$. To define v_1 and v_0 we use a small case distinction. Consider the set C of the children of v_2 that are not spanning the big angle with v_3 at v_2.

1. If C contains a vertex that is not a leaf in $\mathsf{MST}(S)$ or if v_2 has only a single child, we choose it as v_1. We assume w.l.o.g. that v_1 is the successor of v_3 in clockwise order around v_2. . Further, we choose v_0 as a child of v_1 such that v_2 and v_0 are consecutive around v_1 and not spanning the big angle at v_1. If v_1 has two children, and this is true for both, then we choose v_0 such that it is the successor of v_2 in counterclockwise order around v_1. See Figure 2 (a)-(f) for the six different "geometric" (taking the position of the big angles into account) variations of this case. Explicit definitions for these cases can be found below. W.l.o.g., we require the subtree at v_2 to be nonempty in cases (d)-(f) and the subtree of v_1 to be nonempty in cases (c) and (f). For all cases we assume without loss of generality that the angle $v_3v_2v_1$ is clockwise less than π.

   (1a) The edge v_3v_2 is **adjacent** to the big angle at v_2 and the angle $v_2v_1v_0$ is clockwise **less** than π.

   (1b) The edge v_3v_2 is **adjacent** to the big angle at v_2, the angle $v_2v_1v_0$ is clockwise **greater** than π, and the edge v_2v_1 is **adjacent** to the big angle at v_1.

   (1c) The edge v_3v_2 is **adjacent** to the big angle at v_2, the angle $v_2v_1v_0$ is clockwise **greater** than π, and the edge v_2v_1 is **not adjacent** to the big angle at v_1.

   (1d) The edge v_3v_2 is **not adjacent** to the big angle at v_2 and the angle $v_2v_1v_0$ is clockwise **less** than π.

   (1e) The edge v_3v_2 is **not adjacent** to the big angle at v_2, the angle $v_2v_1v_0$ is clockwise **greater** than π, and the edge v_2v_1 is **adjacent** to the big angle at v_1.

   (1f) The edge v_3v_2 is **not adjacent** to the big angle at v_2, the angle $v_2v_1v_0$ is clockwise **greater** than π, and the edge v_2v_1 is **not adjacent** to the big angle at v_1.

2. All vertices in C are leaves in $\mathsf{MST}(S)$. Note that this implies that v_2 has degree exactly three in $\mathsf{MST}(S)$. We choose a vertex of C as v_1 (assuming w.l.o.g. that v_1 is the successor of v_3 in clockwise order around v_2), and choose the other child of v_2 as v_0. Note that if $n \geq 5$ then v_0 cannot be a leaf in $\mathsf{MST}(S)$ and hence v_0 spans a big angle with v_3 at v_2. Hence, taking the location of the big angle at v_2 into account, there are two possibilities for the counterclockwise order of incident edges around v_2. See Figure 3 (a)-(b).

   (2a) The angle $v_1v_2v_0$ is clockwise **less** than π (the edge v_3v_2 is **adjacent** to the big angle at v_2.)

   (2b) The angle $v_1v_2v_0$ is clockwise **greater** than π (the edge v_3v_2 is **not adjacent** to the big angle at v_2.) Here both v_1 and v_0 must be leaves implying that $n = 4$ and we don't have any subtrees.

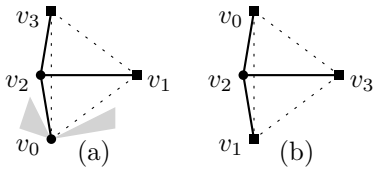

Figure 3: Case 2 for the selection of P and the connections to the rest of $\mathsf{MST}(S)$. Vertices drawn with squares are leaves of $\mathsf{MST}(S)$. The gray triangles indicate possible subtrees of $\mathsf{MST}(S)$ and where they might be connected. Dotted edges are from $\mathsf{MST}^2(P)$. Note that the subtree with root v_0 can be on either side of the supporting line of v_1v_0 and even on both sides as indicated in the figure. Further note that case (b) can appear only if S contains exactly 4 vertices.

The construction of R and B: First we show how to construct the trees $R_P = G(\{v_3, v_2, v_1, v_0\}, E_{PR})$ and $B_P = G(\{v_3, v_2, v_1, v_0\}, E_{PB})$. The vertices of P can either be in convex position or form a triangle with one interior point, with v_1 interior for the cases shown in Figure 2 (b) and (e), and v_2 interior for the cases shown in Figure 2 (e) and (f). Note that there are no other non-convex versions: Otherwise either the path v_3, v_2, v_1, v_0 could not be in $\mathsf{MST}(S)$, or one of the vertices of P could not be incident to a big angle. As in any non-convex case, the complete graph on $\{v_3, v_2, v_1, v_0\}$ is crossing free, any construction of R_P and B_P for the convex case is also valid for the non-convex case. For Case 1 there exist only the two possibilities shown in Figure 4 (a) and (b), as the third possibility for distributing the labels would imply that $\mathsf{MST}(P)$ is self-intersecting, which is a contradiction to the fact that P is a subgraph of $\mathsf{MST}(S)$. In both cases, all edges except the edge v_3v_0 (in E_{PB}) are from $\mathsf{MST}^2(S)$ and have endpoints with different levels in $\mathsf{MST}(S)$ rooted at v_3. In contrast, v_3v_0 is an edge of $\mathsf{MST}^3(S)$, which could be crossed by other edges of the construction. We will later discuss how to handle this.

For Case 2 the coloring for the case shown in Figure 3 (a) is given in Figure 4 (c) and for case shown in Figure 3 (b) the coloring is given in Figure 4 (d). For both cases, all edges are from $\mathsf{MST}^2(S)$.

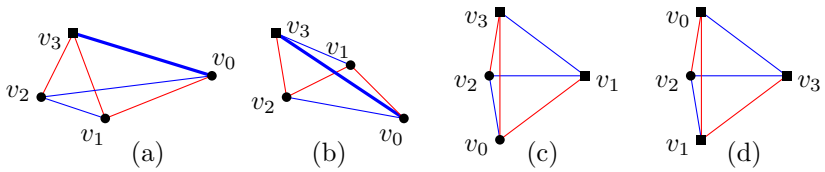

Figure 4: The different colorings for P: (a) for the cases from Figure 2 (a) and (d); (b) for the cases from Figure 2 (b), (c), (e), and (f); (c) for the case from Figure 3 (a); and (d) for the case from Figure 3 (b). The bold edge in (a) and (b) is an edge of $\mathsf{MST}^3(S)$, that is, an edge with link distance 3 in $\mathsf{MST}(S)$.

With R_P and B_P as a base, we now create red and blue trees for all remaining

subtrees of $MST(S)$ and "attach" them to the base. For Case 1 we define three possible subtrees. Let $T'_0 = G(S'_0, E'_0)$ be the subtree (connected component) of $\mathsf{MST}(S)$ that contains v_0 when removing v_1 (and its incident edges) from $\mathsf{MST}(S)$. Likewise, let $T'_1 = G(S'_1, E'_1)$ be the subtree of $\mathsf{MST}(S)$ that contains v_1 when removing v_0 and v_2 from $\mathsf{MST}(S)$, and let $T'_2 = G(S'_2, E'_2)$ be the subtree of $\mathsf{MST}(S)$ that contains v_2 when removing v_1 and v_3 from $\mathsf{MST}(S)$. For Case 2(a) there exists one possible subtree $T'_0 = G(S'_0, E'_0)$, which is the subtree of $\mathsf{MST}(S)$ that contains v_0 when removing v_2 from $\mathsf{MST}(S)$. (Case 2(b) appears only if $n = 4$ and hence the construction is already completed.) The subtrees T'_0, T'_1, and T'_2 are shown as (pairs of) gray triangles in Figure 2 and 3. To connect these subtrees to the bases R_P and B_P, we create corresponding trees T_0, T_1, and T_2 depending on the different shown cases, then apply Construction 1 to them, and possibly recolor them using Lemma 5.

We first consider the different subtrees for Case 1. In essence, for each tree we pick a neighbor from $\{v_3, v_2, v_1, v_0\}$ to add to T'_0, T'_1, T'_2, which we then use as root for Construction 1. When there is a choice we pick a root that is adjacent to the outgoing edge from v_i into the subtree T'_i as defined more precisely below.

$\mathbf{T_0}$: For all cases, we consider the subtree $T_0 = G(S_0, E_0)$ of $\mathsf{MST}(S)$, with $S_0 = S'_0 \cup \{v_1\}$, $E_0 = E'_0 \cup \{v_1v_0\}$. We root T_0 at $r = v_1$ (observe, v_1 is a leaf in T_0 with unique child $s = v_0$) and apply Construction 1 to get $R_0 = G(S_0, E_{0R})$ and $B_0 = G(S_0, E_{0B})$, with the "double-edge" rs removed from both E_{0R} and E_{0B}.

$\mathbf{T_1}$: For the cases depicted in Figure 2(a), (b), (d), and (e), we define $T_1 = G(S_1, E_1)$ of $\mathsf{MST}(S)$, such that $S_1 = S'_1 \cup \{v_0\}$, $E_1 = E'_1 \cup \{v_1v_0\}$. We root T_1 at $r = v_0$ (observe, v_0 is a leaf in T_1 with unique child v_1) and apply Construction 1 to get $R_1 = G(S_1, E_{1R})$ and $B_1 = G(S_1, E_{1B})$, with the "double-edge" rs removed from both E_{1R} and E_{1B}.

In the cases shown in Figure 2(c) and (f), we define $T_1 = G(S_1, E_1)$ of $\mathsf{MST}(S)$, such that $S_1 = S'_1 \cup \{v_2\}$, $E_1 = E'_1 \cup \{v_1v_2\}$. We root T_1 at $r = v_2$ (observe, v_2 is a leaf in T_1 with unique child v_1) and apply Construction 1 to get $R_1 = G(S_1, E_{1R})$ and $B_1 = G(S_1, E_{1B})$, with the "double-edge" rs removed from both E_{1R} and E_{1B}.

$\mathbf{T_2}$: For the cases depicted in Figure 2(a)-(c), let $T_2 = G(S_2, E_2)$ be a subtree of $\mathsf{MST}(S)$, such that $S_2 = S'_2 \cup \{v_1\}$, $E_2 = E'_2 \cup \{v_1v_2\}$. We root T_2 at $r = v_1$ (observe, v_1 is a leaf in T_2 with unique child v_2) and apply Construction 1 to get $R_2 = G(S_2, E_{2R})$ and $B_2 = G(S_2, E_{2B})$, with the "double-edge" rs removed from both E_{2R} and E_{2B}.

For the cases depicted in Figure 2(d)-(f), $T_2 = G(S_2, E_2)$ is the subtree of $\mathsf{MST}(S)$, such that $S_2 = S'_2 \cup \{v_3\}$, $E_2 = E'_2 \cup \{v_3v_2\}$. We root T_2 at $r = v_3$ (observe, v_3 is a leaf in T_2 with unique child v_2) and apply Construction 1 to get $R_2 = G(S_2, E_{2R})$ and $B_2 = G(S_2, E_{2B})$, with the "double-edge" rs removed from both E_{2R} and E_{2B}.

It is easy to see that the edge sets E_{PR}, E_{0R}, E_{1R}, E_{2R}, E_{PB}, E_{0B}, E_{1B}, and E_{2B} are all individually edge disjoint. In the following, we describe how these edge sets are combined to form the two plane spanning trees R and B in

the different cases; see Figure 5 for the convex versions and Figure 6 for the non-convex versions of the cases from Figure 2. For the non-convex cases only points v_1 and v_2 can be in the interior as v_3 or v_0 in the interior would violate Lemma 1. Furthermore, in some of the cases (a)-(f) further restrictions apply as listed below.

(a) Neither v_1 nor v_2 can be the middle point as both clockwise angles $v_3v_2v_1$ and $v_2v_1v_0$ have angle less than π.

(b) Only v_1 can be in the center (the subtree of v_2 is nonempty and hence v_2 wouldn't be incident to a big angle).

(c) Neither v_1 nor v_2 can be in the center (both subtrees are nonempty).

(d) Similar to (a).

(e) Both v_1 and v_2 may be the middle point.

(f) Only v_2 can be in the center (the subtree of v_1 is non-empty).

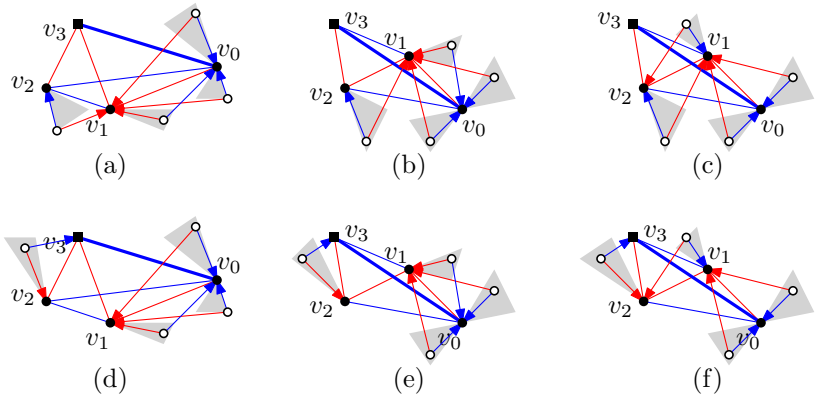

Figure 5: The different plane spanning trees R and B for case 1 when P is in convex position. The bold blue edges v_3v_0 are edges of $\mathsf{MST}^3(S)$, i.e., edges with link distance 3 in $\mathsf{MST}(S)$, that might still be crossed. Note that v_3v_0 can only be crossed in cases (a) and (b), as in the other cases this edge is surrounded by "uncrossable" $\mathsf{MST}^2(S)$-edges.

First we add the red and blue trees for T_0. By construction, only edges of E_{0R} connect to v_1 (only crossing edges of E_{PB}) and the edges of E_{0B} don't cross any edge outside T_0. For the cases (a), (b), (d), and (e) we use the "inverted coloring" (see Lemma 5) for the two trees of T_1. For the remaining cases (c) and (f) we use the "original coloring" (see Lemma 5) for the two trees of T_1. For adding the red and blue trees for T_2 we use the "original coloring" for the cases (a-c), and the "inverted coloring" for the cases (d-f).

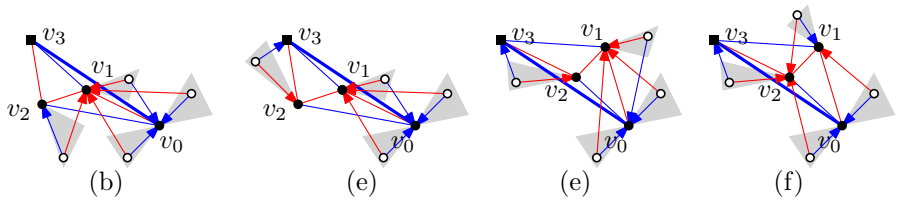

Figure 6: The different plane spanning trees R and B for case 1 when P is not in convex position, with v_1 or v_2 in the interior. The case numberings are the same as the ones in Figure 2.

It can be observed that the resulting graphs $R = G(S, E_{PR} \cup E_{0R} \cup E_{1R} \cup E_{2R})$ and $B = G(S, E_{PB} \cup E_{0B} \cup E_{1B} \cup E_{2B})$ are spanning trees, that $E_R \cap E_B = \emptyset$, and $\max\{\mathsf{BE}(R), \mathsf{BE}(B)\} \leq 3\mathsf{BE}(\mathsf{MST}(S))$ hold, with v_3v_0 from E_{PB} being the only edge (possibly) larger than $2\mathsf{BE}(\mathsf{MST}(S))$.

The $\mathsf{MST}^3(S)$-edge v_3v_0 from E_{PB} is also the only edge that could cause a crossing, see Lemma 3 and Theorem 1. Hence, R is a plane spanning tree. If v_3v_0 is not crossed by any other edge of E_B then also B is a plane spanning tree and we are done. Otherwise, note first that by Lemma 1, the triangles $v_3v_2v_1$ and $v_2v_1v_0$ cannot contain any points of S. Then observe that for the cases in Figure 5 (b), (c), (e) and (f) any edge crossing v_3v_0 that does not have v_0, v_1, v_2 or v_3 as an endpoint must cross an MST-edge between v_0, v_1, v_2 and v_3. This implies that any MST^2-edge that crosses v_3v_0 must have v_0, v_1, v_2 or v_3 as its witness by Lemma 3. By construction, v_0, v_1, v_2 and v_3 are not a witness to any blue edge in the grey subtrees. The edges in E_{PB} are incident to v_0 or v_3 so they also cannot cross v_3v_0.

For cases from Figure 5 (a) and (d), as well as all cases from Figure 6, observe first that v_0v_1, v_1v_2, and v_2v_3 cannot be crossed, again due to Lemma 3 and the fact that v_0, v_1, v_2 and v_3 cannot by a witness to any long edge in the grey subtrees by construction. So any edge crossing with v_3v_0 must have an endpoint in the interior of the convex hull of P or connect directly to v_1 or v_2. The latter however cannot happen: The only points connecting with blue edges to v_1 or v_2 are direct neighbors of these vertices, which reside in the large-angled wedge $v_1v_2v_3$ or $v_0v_1v_2$ respectively. Hence, if v_0v_3 is crossed by some blue edge, there must be a nonempty set $X \subset S \setminus P$ that resides in the interior of the convex hull of P. In the cases depicted in Figure 5 (a) and (d), X lies in the triangle Δ spanned by v_3, v_0, and the intersection of v_3v_1 and v_2v_0. In the cases depicted in Figure 6, the X lies in the triangle Δ spanned by v_3, v_0, and the vertex of P in the interior of the convex hull of P. Further, removing the edge v_0v_3 from B splits B into two connected components that are both plane trees, where v_3 is in one and v_0 is in the other component. Now consider the convex hull of $X \cup \{v_0, v_3\}$, and the path along the boundary of that convex hull between v_0 and v_3 that contains at least one vertex of X. This path contains exactly one edge e that connects the two components of B. Due to the construction of B and R, e can neither be part of R (as the two endpoints of e must reside in two

different subtrees of v_0, v_1 or v_2) nor cross any edge of B (as the only possibly intersected segment of the convex hull boundary of X was v_0v_3). Further, the length of e must be less than $3\mathsf{BE}(\mathsf{MST}(P))$, as e lies inside the triangle Δ, and as all sides of Δ are bounded from above by $3\mathsf{BE}(\mathsf{MST}(P))$. Hence, as v_3v_0 was the only edge that could be crossed by others, the replacement of v_3v_0 by e in B results in two disjoint plane spanning trees R and B with maximum edge length less than $3\mathsf{BE}(\mathsf{MST}(P))$.

As for Case 2 (b), Figure 4 (d) already shows the complete two trees R and B, it remains to consider the subtree for Case 2 (a).

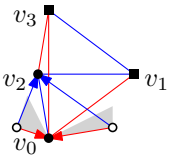

Figure 7: The plane spanning trees R and B for case 2.

$\mathbf{T_0}$: Consider the subtree $T_0 = G(S_0, E_0)$ of $\mathsf{MST}(S)$, with $S_0 = S_0' \cup \{v_2\}$, $E_0 = E_0' \cup \{v_2v_0\}$. We root T_0 at $r = v_2$ (observe, v_2 is a leaf in T_0 with unique child $s = v_0$) and apply Construction 1 to get $R_0 = G(S_0, E_{0R})$ and $B_0 = G(S_0, E_{0B})$, with the "double-edge" rs removed from both E_{0R} and E_{0B}. We use the "inverted coloring" (see Lemma 5) for the two trees of T_0, implying that the edges connecting to v_2 and crossing red edges of E_{PR} are all blue. Hence the graphs $R = G(S, E_{PR} \cup E_{0R})$ and $B = G(S, E_{PB} \cup E_{0B})$ are plane spanning trees, $E_R \cap E_B = \emptyset$, and $\max\{\mathsf{BE}(R), \mathsf{BE}(B)\} \leq 2\mathsf{BE}(\mathsf{MST}(S))$.

This concludes the proof. □

Corollary 5. *For any set S of $n \geq 4$ points in the plane, there exist two plane spanning trees $R = G(S, E_R)$ and $B = G(S, E_B)$ such that $E_R \cap E_B = \emptyset$ and* $\max\{\mathsf{BE}(R), \mathsf{BE}(B)\} \leq 3\mathsf{BE}(\mathsf{MST}(S))$.

We now show that the above construction is worst-case optimal.

Theorem 6. *For any $n > 3$ and $k > 1$ there exists a set S of n points such that for any k disjoint spanning trees, at least one has a bottleneck edge larger than* $(k+1)\mathsf{BE}(\mathsf{MST}(S))$.

Proof. A counterexample simply consists of n points equally distributed on a line segment. (The points can be slightly perturbed to obtain general position.) In this problem instance there are $kn - (k(k+1)/2)$ edges whose distance is strictly less than $(k+1)\mathsf{BE}(\mathsf{MST}(S)) = k+1$. However, we need $kn - k$ edges for k disjoint trees and thus it is impossible to construct that many trees with sufficiently short edges. □

3 Distributed Approach

The previous construction relies heavily on the minimum spanning tree of S. It is well known that this tree cannot be constructed locally, thus we are implicitly assuming that the network is constructed by a single processor that knows the location of all other vertices. In ad-hoc networks, it is often desirable that each vertex can compute its adjacencies using only local information.

In the following, we provide an alternative construction. Although the length of the edges is increased by a constant factor, it has the benefit that it can be constructed locally and that it can be extended to compute k layers. The only global property that is needed is a value β that should be at least $\mathsf{BE}(\mathsf{MST}(S))$. We also note that these plane disjoint graphs are not necessarily trees, as large cycles cannot be detected locally.

Before we describe our approach, we report the result of García [5] that states that every point set of at least $3k$ points contains k layers. Since the details of this construction are important for our construction and the manuscript is not yet available, we add a proof sketch.

Theorem 7 ([5])**.** *Every point set that consists of at least* $3k$ *points contains* k *layers.*

Proof. First, recall that for every set of n points, there is a *center point* c such that every line through c splits the point set into two parts that each contain at least $\lfloor n/3 \rfloor$ points. For ease of explanation, we assume that every line through c contains at most one point. Number the points $v_0, v_1, \ldots, v_{n-1}$ in clockwise circular order around c. We split the plane into three angular regions by the three rays originating from c and passing through v_0, $v_{\lfloor \frac{n}{3} \rfloor}$, and $v_{\lfloor \frac{2n}{3} \rfloor}$, see Figure 8. Since every line through the center contains at least $n/3$ points on each side, the three angular regions are convex. We declare v_0 to be the representative of the angular region between the rays through v_0 and $v_{\lfloor \frac{n}{3} \rfloor}$ and connect the vertices $v_1, ..., v_{\lfloor \frac{n}{3} \rfloor}$ in this region to v_0. Similarly, we assign $v_{\lfloor \frac{n}{3} \rfloor}$ to be the representative of angle between the rays center through $v_{\lfloor \frac{n}{3} \rfloor}$ and $v_{\lfloor \frac{2n}{3} \rfloor}$ and connect vertices $v_{\lfloor \frac{n}{3} \rfloor+1}, ..., v_{\lfloor \frac{2n}{3} \rfloor}$ to $v_{\lfloor \frac{n}{3} \rfloor}$. Finally, we connect vertices $v_{\lfloor \frac{2n}{3} \rfloor+1}, ..., v_{n-1}$ to $v_{\lfloor \frac{2n}{3} \rfloor}$. This results in a non-crossing spanning tree.

For the second tree, we rotate the construction and we use v_1, $v_{\lfloor \frac{n}{3} \rfloor+1}$, and $v_{\lfloor \frac{2n}{3} \rfloor+1}$ to define the three regions, and so on. □

While this construction provides a simple method of constructing the k layers, it does not give any guarantee on the length of the longest edge in this construction. To give such a guarantee, we combine it with a bucketing approach: we partition the point set using a grid (whose size will depend on k and β), solve the problem in each box with sufficiently many points independently, and then combine the subproblems to obtain a global solution (see Figure 9).

We place a grid with cells of height and width $6k\beta$ and classify the points according to which grid cell contains them (if a point lies exactly on the separating lines, pick an arbitrary adjacent cell). We say that a grid cell is a *dense box* if it contains at least $3k$ points of S. Similarly, it is a *sparse box* if it contains

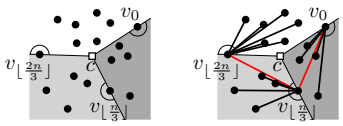

Figure 8: Extracting one layer: (a) The three sectors defined by v_0, $v_{\lfloor \frac{n}{3} \rfloor}$, and $v_{\lfloor \frac{2n}{3} \rfloor}$. (b) Connecting the points to the representative of their sector. The red edges connect the representatives.

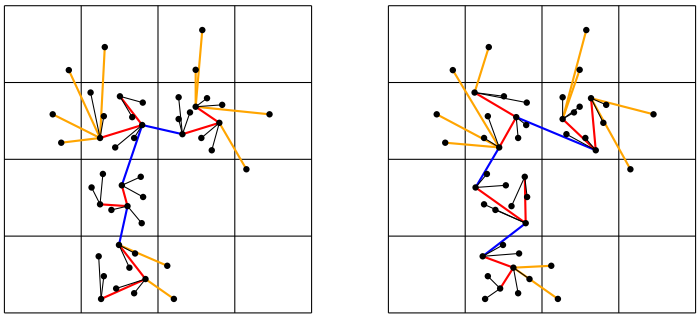

Figure 9: The distributed approach: a grid is placed over the point set and different representatives construct different graphs ((a) and (b)). The red and black edges form the tree in each dense cell, blue edges connect the dense cells, and orange edges connect the vertices in sparse cells.

points of S but is not dense. We observe that dense and sparse boxes satisfy the following properties.

Lemma 6. *Given two non-adjacent boxes B and B', the points in B and B' cannot be connected by edges of length at most β using only points from sparse boxes.*

Proof. Suppose the contrary and let B and B' be two boxes s.t. there is a path that uses edges of length at most β between a point in B to a point in B' visiting only points in sparse boxes. This path crosses the sides of a certain number of boxes in a given order; let σ be the sequence of these sides, with adjacent duplicates removed. Observe first that horizontal and vertical sides alternate in σ, as otherwise the path would have to use at least $6k-1$ points to traverse a sparse box, but there are only at most $3k-1$. Since B and B' are non-adjacent, w.l.o.g., there is a vertical side s that has two adjacent horizontal sides in σ with different y-coordinates. Hence, between the two horizontal sides, the corresponding part of the path has length at least $6k\beta$, and may use only the points in the two boxes adjacent to s. But since any sparse box contains at most $3k-1$ points and the distance between two consecutive points along the

path is at most β, that part of the path can have length at most $(6k-1)\beta$, a contradiction. □

Corollary 8. *Dense boxes are connected by the 8-neighbor topology.*

Lemma 7. *Any set S of at least $4 \cdot (3k-1)+1$ points with $\beta \geq \mathsf{BE}(\mathsf{MST}(S))$ contains at least one dense box.*

Proof. Assume S consists of only sparse boxes. This implies that the points are distributed over at least five boxes, and thus, there is a pair of boxes that is non-adjacent. Using Lemma 6, this means that these boxes cannot be connected using edges of length at most $\mathsf{BE}(\mathsf{MST}(S))$, a contradiction. □

Lemma 8. *In any set S of at least $4 \cdot (3k-1)+1$ points with $\beta \geq \mathsf{BE}(\mathsf{MST}(S))$, all sparse boxes are adjacent to a dense box.*

Proof. This follows from Lemma 6, since any sparse box that is not adjacent to a dense box cannot be connected to any dense box using edges of length at most $\beta \geq \mathsf{BE}(\mathsf{MST}(S))$. □

Next, we assign all points to dense boxes. In order to do this, let c_B be the center of a dense box B. Note that c_B is not necessarily the center point of the points in this box. We consider the Voronoi diagram of the centers of all dense boxes and assign a point p to B if p lies in the Voronoi cell of c_B. Let S_B be the set of points of S that are associated with a dense box B. We note that each dense box B gets assigned at least all points in its own box, since in the case of adjacent dense boxes, the boundary of the Voronoi cell coincides with the shared boundary of these boxes (see Figure 10).

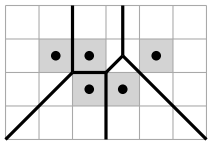

Figure 10: The Voronoi cells of the centers of the dense boxes.

Furthermore, we can compute the points assigned to each box locally. By Lemma 8 all sparse boxes are adjacent to a dense box, and hence for any point p in a sparse box B its distance to its nearest center is at most $3\ell/\sqrt{2}$, where $\ell = 6k\beta$. It follows that only the centers of cells of neighbors and neighbors of neighbors need to be considered.

Lemma 9. *For any two dense boxes B and B', we have that the convex hulls of S_B and $S_{B'}$ are disjoint.*

Proof. We observe that the convex hull of S_B is contained in the Voronoi cell of c_B. Hence, since the Voronoi cells of different dense boxes are disjoint, the convex hulls of the points assigned to them are also disjoint. □

For each dense box B, we apply Theorem 7 on the points inside the dense box to compute k disjoint layers of S_B. Next, we connect all sparse points in S_B to the representative of the sector that contains them in each layer. Since all points in the same sector connect to the same representative and the sectors of the same layer do not overlap, we obtain a plane graph for each layer within the convex hull of each S_B.

Hence, we obtain k pairwise disjoint layers such that in each layer the points associated to each dense box are connected. Moreover, since the created edges stay within the convex hull of each subproblem and by Lemma 9 those hulls are disjoint, each layer is plane. Thus, to assure that each layer is connected, we must connect the construction between dense boxes.

We connect adjacent dense boxes in a tree-like manner using the following rules:

- Always connect a dense box to the dense box below it.
- Always connect a dense box to the dense box to the left of it.
- If neither the box below nor the one to the left of it is dense, connect the box to the dense box diagonally below and to the left of it.
- If neither the box above nor the one to the left of it is dense, connect the box to the dense box diagonally above and to the left of it.

To connect two dense boxes, we find and connect two representatives p and q (one from each dense box) such that p lies in the sector of q and q lies in the sector of p; see Figure 11 (a).

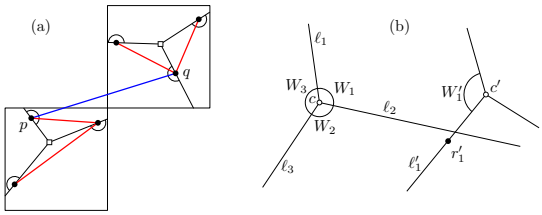

Figure 11: Connecting two dense boxes by means of p and q. The half-circles in (a) indicate which sector each representative covers. The red edges connect the dense boxes internally and the blue edge connects the two dense cells. (b) illustrates the sectors involved in connecting two neighboring dense boxes.

Lemma 10. *For any layer and any two adjacent dense boxes B and B', there are two representatives p and q in B and B', respectively, such that p lies in the sector of q and q lies in the sector of p.*

Proof. Consider two boxes B and B' with center points (of their respective point sets) c and c'. Now let W_1 and W_1' with representatives r_1 and r_1' denote the sectors containing c' and c, respectively; see Figure 11. The other sectors W_2

and W_3 of B with representatives r_2 and r_3 are ordered clockwise. We use ℓ_i to denote the ray from c containing r_i. If $r_1 \in W_1'$ and $r_1' \in W_1$ we are done. So assume that $r_1' \notin W_1$, the case when $r_1 \notin W_1'$ (or when both $r_1 \notin W_1'$ and $r_1' \notin W_1$) is symmetric. It follows that r_1' is in sector W_2 if the line segment $c'r_1'$ intersects ℓ_2 or sector W_3 if the segment intersects ℓ_2 and ℓ_3. Assume that r_1' is in sector W_2 (again the argument is symmetric when r_1' is in sector W_3). Now r_2 can be positioned on ℓ_2 between c and the intersection point with $c'r_1'$ or behind this intersection point when viewed from c. In the former case r_1' is in W_2 and r_2 is in W_1' and we are done. In the latter case the segments cr_2 and $c'r_1'$ cross. Since $c, r_2 \in B$ and $c', r_1' \in B'$ this crossing would imply that B and B' are not disjoint, a contradiction. □

Now that we have completed the description of the construction, we show that each layer of the resulting graph is plane and connected, and that the length of the longest edge is bounded.

Lemma 11. *Each layer is plane.*

Proof. Since dense boxes are internally plane and the addition of edges to the sparse points do not violate planarity, it suffices to show that the edges between dense boxes cannot cross any previously inserted edges and that these edges cannot intersect other edges used to connect dense boxes.

We first show that the edge used to connect boxes B and B' is contained in the union of the Voronoi cells of these two boxes. If B and B' are horizontally or vertically adjacent, the connecting edge stays in the union of the two dense boxes, which is contained in their Voronoi cells. If B and B' are diagonally adjacent, we connect them only if their shared horizontal and vertical neighbors are not dense. This implies that at least the two triangles defined by the sides of B and B' that are adjacent to their contact point are part of the union of the Voronoi cells of these boxes. Hence, the edge used to connect B and B' cannot intersect the Voronoi cell of any other dense box. Since all points of a dense box in a sector connect to the same representative and these edges lie entirely inside the sector, the edge connecting two adjacent boxes can intersect only at one of the two representatives, but does not cross them. Therefore, an edge connecting two adjacent dense boxes by connecting the corresponding representatives cannot cross any previously inserted edge.

Next, we show that edges connecting two dense boxes cannot cross. Since any edge connecting two dense boxes stays within the union of the Voronoi cells of B and B', the only way for two edges to intersect is if they connect to the same box B and intersect in the Voronoi cell of B. If the connecting edges lie in the same sector of B, they connect to the same representative and thus they cannot cross. If they lie in different sectors of B, the edges lie entirely inside their respective sectors. Since these sectors are disjoint, this implies that the edges cannot intersect. □

Lemma 12. *Each layer is connected.*

Proof. Since the sectors of the representatives of the dense boxes cover the plane, each point in a sparse box is connected to a representative of the dense box it is assigned to. Hence, showing that the dense boxes are connected, completes the proof.

By Corollary 8, the dense boxes are connected using the 8-neighbor topology. This implies that there is a path between any pair of dense boxes where every step is one to a horizontally, vertically, or diagonally adjacent box. Since we always connect horizontally or vertically adjacent boxes and we connect diagonally adjacent boxes when they share no horizontal and vertical dense neighbor, the layer is connected after adding edges as described in the proof of Lemma 10. □

Lemma 13. *The distance between a representative in a dense box B and any point connecting to it is at most $12\sqrt{2}k\beta$.*

Proof. Since the representatives of B are connected only to points from dense and sparse boxes adjacent B, the distance between a representative and a point connected to it is at most the length of the diagonal of the 2×2 grid cells with B as one of its boxes. Since a box has width $6k\beta$, this diagonal has length $2\sqrt{2} \cdot 6k\beta = 12\sqrt{2}k\beta$. □

Theorem 9. *For all point sets with at least $4(3k-1)+1$ points, we can extract k plane layers with the longest edge having length at most $12\sqrt{2}k\mathsf{BE}(\mathsf{MST}(S))$.*

Acknowledgements

We would like to thank all participants of the 10th European Research Week on Geometric Graphs for fruitful discussions. O.A., A.P., and B.V. were partially supported by the ESF EUROCORES programme EuroGIGA - ComPoSe, Austrian Science Fund (FWF): I 648-N18. T.H. was supported by the Austrian Science Fund (FWF): P23629-N18 'Combinatorial Problems on Geometric Graphs'. M.K. was supported in part by the ELC project (MEXT KAKENHI No. 12H00855 and 15H02665). A.P. is supported by an Erwin Schrödinger fellowship, Austrian Science Fund (FWF): J-3847-N35.

References

[1] Aichholzer, O., Hackl, T., Korman, M., van Kreveld, M.J., Löffler, M., Pilz, A., Speckmann, B., Welzl, E.: Packing plane spanning trees and paths in complete geometric graphs. In: Proc. Canadian Conference on Computational Geometry (CCCG). pp. 233–238 (2014)

[2] Bose, P., Morin, P., Stojmenovic, I., Urrutia, J.: Routing with guaranteed delivery in ad hoc wireless networks. Wireless Networks 7(6), 609–616 (2001)

[3] Dor, D., Tarsi, M.: Graph decomposition is NP-complete: A complete proof of Holyer's conjecture. SIAM J. Comput. 26(4), 1166–1187 (1997)

[4] Fussen, M., Wattenhofer, R., Zollinger, A.: Interference arises at the receiver. In: Proc. International Conference on Wireless Networks, Communications, and Mobile Computing (WIRELESSCOM). pp. 427–432 (2005)

[5] García, A.: On the number of edge-disjoint plane spanning trees of K_n (2015), unpublished manuscript

[6] Korman, M.: Minimizing interference in ad-hoc networks with bounded communication radius. Information Processing Letters 112(19), 748–752 (2012)

[7] Kranakis, E., Singh, H., Urrutia, J.: Compass routing on geometric networks. In: Proc. Canadian Conference on Computational Geometry (CCCG). pp. 51–54 (1999)

[8] Priesler, M., Tarsi, M.: Multigraph decomposition into stars and into multistars. Discrete Mathematics 296(2-3), 235–244 (2005)

[9] Tarsi, M.: Decomposition of a complete multigraph into simple paths: Nonbalanced handcuffed designs. J. Comb. Theory, Ser. A 34(1), 60–70 (1983)